

7-31-00

A

Please type a plus sign (+) inside this box [+]

PTO/SB/05 (12/97)

Approved for use through 09/30/00. OMB 0651-0032

Patent and Trademark Office: U.S. DEPARTMENT OF COMMERCE

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

**UTILITY PATENT APPLICATION TRANSMITTAL**  
(Only for new nonprovisional applications under 37 CFR 1.53(b))Attorney Docket No. 42390.P7876XTotal Pages 2First Named Inventor or Application Identifier Hooper, et al.Express Mail Label No. EL234218183USJC857 U.S. PRO  
09/626535  
07/27/00ADDRESS TO: Assistant Commissioner for Patents  
Box Patent Application  
Washington, D. C. 20231**APPLICATION ELEMENTS**

See MPEP chapter 600 concerning utility patent application contents.

1. X Fee Transmittal Form  
(Submit an original, and a duplicate for fee processing)
2. X Specification (Total Pages 31)  
(preferred arrangement set forth below)
  - Descriptive Title of the Invention
  - Cross References to Related Applications
  - Statement Regarding Fed sponsored R & D
  - Reference to Microfiche Appendix
  - Background of the Invention
  - Brief Summary of the Invention
  - Brief Description of the Drawings (if filed)
  - Detailed Description
  - Claims
  - Abstract of the Disclosure
3. X Drawings(s) (35 USC 113) (Total Sheets 15)
4. X Oath or Declaration (Total Pages 5) (Unsigned)
  - a.      Newly Executed (Original or Copy)
  - b.      Copy from a Prior Application (37 CFR 1.63(d))  
(for Continuation/Divisional with Box 17 completed) (**Note Box 5 below**)
  - i.      DELETIONS OF INVENTOR(S) Signed statement attached deleting inventor(s) named in the prior application, see 37 CFR 1.63(d)(2) and 1.33(b).
5.      Incorporation By Reference (useable if Box 4b is checked)  
The entire disclosure of the prior application, from which a copy of the oath or declaration is supplied under Box 4b, is considered as being part of the disclosure of the accompanying application and is hereby incorporated by reference therein.
6.      Microfiche Computer Program (Appendix)

JC813 U.S. PRO  
07/27/00

09626535 "072700"

7. ☐ Nucleotide and/or Amino Acid Sequence Submission  
(if applicable, all necessary)  
a. ☐ Computer Readable Copy  
b. ☐ Paper Copy (identical to computer copy)  
c. ☐ Statement verifying identity of above copies

**ACCOMPANYING APPLICATION PARTS**

8. ☐ Assignment Papers (cover sheet & documents(s))  
9. ☐ a. 37 CFR 3.73(b) Statement (where there is an assignee)  
☐ b. Power of Attorney  
10. ☐ English Translation Document (if applicable)  
11. ☐ a. Information Disclosure Statement (IDS)/PTO-1449  
☐ b. Copies of IDS Citations  
12. ☐ Preliminary Amendment  
13. ☒ Return Receipt Postcard (MPEP 503) (Should be specifically itemized)  
14. ☐ a. Small Entity Statement(s)  
☐ b. Statement filed in prior application, Status still proper and desired  
15. ☐ Certified Copy of Priority Document(s) (if foreign priority is claimed)  
16. ☐ Other:

17. **If a CONTINUING APPLICATION**, check appropriate box and supply the requisite information:  
☐ Continuation ☐ Divisional ☒ Continuation-in-part (CIP)  
of prior application No: 09/475,614, filed 12/30/99.

18. **Correspondence Address**

☐ Customer Number or Bar Code Label   
(Insert Customer No. or Attach Bar Code Label here)  
or  
☒ Correspondence Address Below

NAME Charles K. Young, Reg. #39,435 Charles K. Young  
BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN LLP

ADDRESS 12400 Wilshire Boulevard  
Seventh Floor

CITY Los Angeles STATE California ZIP CODE 90025-1026

Country U.S.A. TELEPHONE (408) 720-8598 FAX (408) 720-9397



# MULTI-THREADED SCHEDULED RECEIVE FOR FAST NETWORK PORT DATA

5

## BACKGROUND OF THE INVENTION

### 1. Field of the Invention

The described invention relates to the field of network communications. In particular, the invention relates to a method for scheduling multiple threads to process  
10 incoming network data.

### 2. Description of Related Art

Networking products such as routers require high speed components for packet  
15 data movement, i.e., collecting packet data from incoming network device ports and queuing the packet data for transfer to appropriate forwarding device ports. They also require high-speed special controllers for processing the packet data, that is, parsing the data and making forwarding decisions. Because the implementation of these high-speed functions usually involves the development of ASIC or custom devices, such  
20 networking products are of limited flexibility. For example, each controller is assigned to service network packets from one or more given ports on a permanent basis.

## BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a block diagram of a communication system employing a hardware-based multi-threaded processor.

5           FIG. 2 is a block diagram of a microengine employed in the hardware-based multi-threaded processor of FIG. 1.

FIG. 3 is an illustration of an exemplary thread task assignment.

FIG. 4 is a block diagram of an I/O bus interface shown in FIG. 1.

FIG. 5 is a detailed diagram of a bus interface unit employed by the I/O bus  
10   interface of FIG. 4.

FIGS. 6A-6F are illustrations of various bus configuration control and status registers (CSRs).

FIG. 7A is a detailed diagram illustrating the interconnection between a plurality of 10/100 Ethernet (“slow”) ports and the bus interface unit.

15           FIG. 7B is a detailed diagram illustrating the interconnection between two Gigabit Ethernet (“fast”) ports and the bus interface unit.

FIGS. 8A-8C are illustrations of the formats of the RCV\_RDY\_CTL, RCV\_RDY\_HI and RCV\_RDY\_LO CSR registers, respectively.

FIG. 9 is a depiction of the receive threads and their interaction with the I/O  
20   bus interface during a receive process.

FIGS. 10A and 10B are illustrations of the format of the RCV\_REQ FIFO and the RCV\_CTL FIFO, respectively.

FIG. 11 is an illustration of the thread done registers.

FIG. 12 shows a graphical overview of the scheduling process showing an  
25   exemplary receive scheduler and three receive threads.



## DETAILED DESCRIPTION

Referring to FIG. 1, a communication system 10 includes a parallel,  
hardware-based multi-threaded processor 12. The hardware based multi-threaded  
processor 12 is coupled to a first peripheral bus (shown as a PCI bus) 14, a second  
peripheral bus referred to as an I/O bus 16 and a memory system 18. The system 10 is  
especially useful for tasks that can be broken into parallel subtasks or functions. The  
hardware-based multi-threaded processor 12 includes multiple microengines 22, each  
with multiple hardware controlled program threads that can be simultaneously active  
and independently work on a task. In the embodiment shown, there are six  
microengines 22a-22f and each of the six microengines is capable of processing four  
program threads, as will be described more fully below.

The hardware-based multi-threaded processor 12 also includes a processor  
23 that assists in loading microcode control for other resources of the hardware-based  
multi-threaded processor 12 and performs other general purpose computer type  
functions such as handling protocols, exceptions, extra support for packet processing  
where the microengines pass the packets off for more detailed processing. In one  
embodiment, the processor 23 is a StrongARM (ARM is a trademark of ARM Limited,  
United Kingdom) core based architecture. The processor (or core) 23 has an operating  
system through which the processor 23 can call functions to operate on the  
microengines 22a-22f. The processor 23 can use any supported operating system,  
preferably real-time operating system. For the core processor implemented as a  
StrongARM architecture, operating systems such as MicrosoftNT real-time, VXWorks  
and :CUS, a freeware operating system available over the Internet, can be used.

The six microengines 22a-22f each operate with shared resources including  
the memory system 18, a PCI bus interface 24 and an I/O bus interface 28. The PCI bus

interface provides an interface to the PCI bus 14. The I/O bus interface 28 is responsible for controlling and interfacing the processor 12 to the I/O bus 16. The memory system 18 includes a Synchronous Dynamic Random Access Memory (SDRAM) 18a, which is accessed via an SDRAM controller 26a, a Static Random Access Memory (SRAM) 18b, which is accessed using an SRAM controller 26b, and a nonvolatile memory (shown as a FlashROM) 18c that is used for boot operations. The SDRAM 16a and SDRAM controller 26a are typically used for processing large volumes of data, e.g., processing of payloads from network packets. The SRAM 18b and SRAM controller 26b are used in a networking implementation for low latency, fast access tasks, e.g., accessing look-up tables, memory for the processor 23, and so forth. The microengines 22a-22f can execute memory reference instructions to either the SDRAM controller 26a or the SRAM controller 18b.

The hardware-based multi-threaded processor 12 interfaces to network devices such as a media access controller device, including a “slow” device 30 (e.g., 10/100BaseT Ethernet MAC) and/or a “fast” device 31, such as Gigabit Ethernet MAC, ATM device or the like, over the I/O Bus 16. In the embodiment shown, the slow device 30 is an 10/100 BaseT Octal MAC device and thus includes 8 slow ports 32a-32h, and the fast device is a Dual Gigabit MAC device having two fast ports 33a, 33b. Each of the network devices attached to the I/O Bus 16 can include a plurality of ports to be serviced by the processor 12. Other devices, such as a host computer (not shown), that may be coupled to the PCI bus 14 are also serviced by the processor 12. In general, as a network processor, the processor 12 can interface to any type of communication device or interface that receives/sends large amounts of data. The processor 12 functioning as a network processor could receive units of packet data from the devices 30, 31 and process those units of packet data in a parallel manner, as will be described.

The unit of packet data could include an entire network packet (e.g., Ethernet packet) or a portion of such a packet.

Each of the functional units of the processor 12 are coupled to one or more internal buses. The internal buses include an internal core bus 34 (labeled “AMBA”) for coupling the processor 23 to the memory controllers 26a, 26b and to an AMBA translator 36. The processor 12 also includes a private bus 38 that couples the microengines 22a-22f to the SRAM controller 26b, AMBA translator 36 and the Fbus interface 28. A memory bus 40 couples the memory controllers 26a, 26b to the bus interfaces 24, 28 and the memory system 18.

Referring to FIG. 3, an exemplary one of the microengines 22a-22f is shown. The microengine 22a includes a control store 70 for storing a microprogram. The microprogram is loadable by the central processor 20. The microengine 70 also includes control logic 72. The control logic 72 includes an instruction decoder 73 and program counter units 72a-72d. The four program counters are maintained in hardware. The microengine 22a also includes context event switching logic 74. The context event switching logic 74 receives messages (e.g., SEQ\_#\_EVENT\_RESPONSE; FBI\_EVENT\_RESPONSE; SRAM\_EVENT\_RESPONSE; SDRAM\_EVENT\_RESPONSE; and AMBA\_EVENT\_RESPONSE) from each one of the share resources, e.g., SRAM 26b, SDRAM 26a, or processor core 20, control and status registers, and so forth. These messages provides information on whether a requested function has completed. Based on whether or not the function requested by a thread has completed and signaled completion, the thread needs to wait for that complete signal, and if the thread is enabled to operate, then the thread is placed on an available thread list (not shown). As earlier mentioned, in one embodiment, the microengine 22a can have a maximum of four threads of execution available.

In addition to event signals that are local to an executing thread, the microengine employs signaling states that are global. With signaling states, an executing thread can broadcast a signal state to all microengines 22. Any and all threads in the microengines can branch on these signaling states. These signaling states can be used to determine availability of a resource or whether a resource is due for servicing.

The context event logic 74 has arbitration for the four threads. In one embodiment, the arbitration is a round robin mechanism. However, other arbitration techniques, such as priority queuing or weighted fair queuing, could be used. The microengine 22a also includes an execution box (EBOX) data path 76 that includes an arithmetic logic unit (ALU) 76a and a general purpose register (GPR) set 76b. The ALU 76a performs arithmetic and logical functions as well as shift functions.

The microengine 22a further includes a write transfer register file 78 and a read transfer register file 80. The write transfer register file 78 stores data to be written to a resource. The read transfer register file 80 is for storing return data from a resource. Subsequent to or concurrent with the data arrival, an event signal from the respective shared resource, e.g., memory controllers 26a, 26b, or core 23, will be provided to the context event arbiter 74, which in turn alerts the thread that the data is available or has been sent. Both transfer register files 78, 80 are connected to the EBOX 76 through a data path. In the described implementation, each of the register files includes 64 registers.

The functionality of the microengine threads is determined by microcode loaded (via the core processor) for a particular user's application into each microengine's control store 70. Referring to FIG. 3, an exemplary thread task assignment 90 is shown. Typically, one of the microengine threads is assigned to serve as a receive scheduler 92 and another as a transmit scheduler 94. A plurality of threads

are configured as receive processing threads 96 and transmit processing (or “fill”) threads 98. Other thread task assignments include a transmit arbiter 100 and one or more core communication threads 102. Once launched, a thread performs its function independently.

5           The receive scheduler thread 92 assigns packets to receive processing threads 96. In a packet forwarding application for a bridge/router, for example, the receive processing thread parses packet headers and performs lookups based in the packet header information. Once the receive processing thread or threads 96 has processed the packet, it either sends the packet as an exception to be further processed  
10   by the core 23 (e.g., the forwarding information cannot be located in lookup and the core processor must learn it), or stores the packet in the SDRAM and queues the packet in a transmit queue by placing a packet link descriptor for it in a transmit queue associated with the transmit (forwarding port) indicated by the header/lookup. The transmit queue is stored in the SRAM. The transmit arbiter thread 100 prioritizes the  
15   transmit queues and the transmit scheduler thread 94 assigns packets to transmit processing threads that send the packet out onto the forwarding port indicated by the header/lookup information during the receive processing.

          The receive processing threads 96 may be dedicated to servicing particular ports or may be assigned to ports dynamically by the receive scheduler thread 92. For  
20   certain system configurations, a dedicated assignment may be desirable. For example, if the number of ports is equal to the number of receive processing threads 96, then it may be quite practical as well as efficient to assign the receive processing threads to ports in a one-to-one, dedicated assignment. In other system configurations, a dynamic assignment may provide a more efficient use of system resources.

25           The receive scheduler thread 92 maintains scheduling information 104 in the GPRs 76b of the microengine within which it executes. The scheduling information

104 includes thread capabilities information 106, port-to-thread assignments (list) 108 and “thread busy” tracking information 110. At minimum, the thread capabilities information informs the receive scheduler thread as to the type of tasks for which the other threads are configured, e.g., which threads serve as receive processing threads.

5 Additionally, it may inform the receive scheduler of other capabilities that may be appropriate to the servicing of a particular port. For instance, a receive processing thread may be configured to support a certain protocol, or a particular port or ports. A current list of the ports to which active receive processing threads have been assigned by the receive scheduler thread is maintained in the thread-to-port assignments list 108.

10 The thread busy mask register 110 indicates which threads are actively servicing a port. The receive scheduler uses all of this scheduling information in selecting threads to be assigned to ports that require service for available packet data, as will be described in further detail below.

Referring to FIG. 4, the I/O bus interface 28 includes shared resources 120,  
15 which are coupled to a push/pull engine interface 122 and a bus interface unit 124. The bus interface unit 124 includes a ready bus controller 126 connected to a ready bus 128 and an Fbus controller 130 for connecting to a portion of the I/O bus referred to as an Fbus 132. Collectively, the ready bus 128 and the Fbus 132 make up the signals of the I/O bus 16 (FIG. 1). The resources 120 include two FIFOs, a transmit FIFO 134 and a  
20 receive FIFO 136, as well as CSRs 138, a scratchpad memory 140 and a hash unit 142. The Fbus 132 transfers data between the ports of the devices 30, 31 and the I/O bus interface 28. The ready bus 128 is an 8-bit bus that performs several functions. It is used to read control information about data availability from the devices 30, 31, e.g., in the form of ready status flags. It also provides flow control information to the devices  
25 30, 31, and may be used to communicate with another network processor 12 that is connected to the Fbus 132. Both buses 128, 132 are accessed by the microengines 22

through the CSRs 138. The CSRs 138 are used for bus configuration, for accessing the bus interface unit 124, and for inter-thread signaling. They also include a several counters and thread status registers, as will be described. The CSRs 138 are accessed by the microengines 22 and the core 23. The receive FIFO (RFIFO) 136 includes data buffers for holding data received from the Fbus 132 and is read by the microengines 22. The transmit FIFO (TFIFO) 134 includes data buffers that hold data to be transmitted to the Fbus 132 and is written by the microengines 22. The scatchpad memory 140 is accessed by the core 23 and microengines 22, and supports a variety of operations, including read and write operations, as well as bit test, bit test/clear and increment operations. The hash unit 142 generates hash indexes for 48-bit or 64-bit data and is accessed by the microengines 22 during lookup operations.

The processors 23 and 22 issue commands to the push/pull engine interface 122 when accessing one of the resources 120. The push/pull engine interface 122 places the commands into queues (not shown), arbitrates which commands to service, and moves data between the resources 120, the core 23 and the microengines 22. In addition to servicing requests from the core 23 and microengines 22, the push/pull engines 122 also service requests from the ready bus 128 to transfer control information to a register in the microengine read transfer registers 80.

When a thread issues a request to a resource 120, a command is driven onto an internal command bus 150 and placed in queues within the push/pull engine interface 122. Receive/read-related instructions (such as instructions for reading the CSRS) are written to a “push” command queue.

The CSRs 138 include the following types of registers: Fbus receive and transmit registers; Fbus and ready bus configuration registers; ready bus control registers; hash unit configuration registers; interrupt registers; and several miscellaneous

registers, including a thread status registers. Those of the registers which pertain to the receive process will be described in further detail.

The interrupt/signal registers include an INTER\_THD\_SIG register for inter-thread signaling. Any thread within the microengines 22 or the core 23 can write a thread number to this register to signal an inter-thread event.

Further details of the Fbus controller 130 and the ready bus controller 126 are shown in FIG. 5. The ready bus controller 126 includes a programmable sequencer 160 for retrieving MAC device status information from the MAC devices 30, 31, and asserting flow control to the MAC devices over the ready bus 128 via ready bus interface logic 161. The Fbus controller 130 includes Fbus interface logic 162, which is used to transfer data to and from the devices 30, 31, is controlled by a transmit state machine (TSM) 164 and a receive state machine (RSM) 166. In the embodiment herein, the Fbus 132 may be configured as a bidirectional 64-bit bus, or two dedicated 32-bit buses. In the unidirectional, 32-bit configuration, each of the state machines owns its own 32-bit bus. In the bidirectional configuration, the ownership of the bus is established through arbitration. Accordingly, the Fbus controller 130 further includes a bus arbiter 168 for selecting which state machine owns the Fbus 132.

Some of the relevant CSRs used to program and control the ready bus 128 and Fbus 132 for receive processes are shown in FIGS. 6A-6F. Referring to FIG. 6A, RDYBUS\_TEMPLATE\_PROGx registers 170 are used to store instructions for the ready bus sequencer. Each register of these 32-bit registers 170a, 170b, 170c, includes four, 8-bit instruction fields 172. Referring to FIG. 6B, a RCV\_RDY\_CTL register 174 specifies the behavior of the receive state machine 166. The format is as follows: a reserved field (bits 31:15) 174a; a fast port mode field (bits 14:13) 174b, which specifies the fast (Gigabit) port thread mode, as will be described; an auto push prevent window field (bits 12:10) 174c for specifying the autopush prevent window used by the

ready bus sequencer to prevent the receive scheduler from accessing its read transfer registers when an autopush operation (which pushes information to those registers) is about to begin; an autopush enable (bit 9) 174d, used to enable autopush of the receive ready flags; another reserved field (bit 8) 174e; an autopush destination field (bits 7:6) 174f for specifying an autopush operation's destination register; a signal thread enable field (bit 5) 174g which, when set, indicates the thread to be signaled after an autopush operation; and a receive scheduler thread ID (bits 4:0) 174h, which specifies the ID of the microengine thread that has been configured as a receive scheduler.

Referring to FIG. 6C, a REC\_FASTPORT\_CTL register 176 is relevant to receiving packet data from fast ports (fast port mode) only. It enables receive threads to view the current assignment of header and body thread assignments for the two fast ports, as will be described. It includes the following fields: a reserved field (bits 31:20) 176a; an FP2\_HDR\_THD\_ID field (bits 19:15) 176b, which specifies the fast port 2 header receive (processing) thread ID; an FP2\_BODY\_THD\_ID field (bits 14:10) 176c for specifying the fast port 2 body receive processing thread ID; an FP1\_HDR\_THD\_ID field (bits 9:5) 176d for specifying the fast port 1 header receive processing thread ID; and an FP1\_BODY\_THD\_ID field (bits 4:0) 176e for specifying the fast port 1 body processing thread ID. The manner in which these fields are used by the RSM 166 will be described in detail later.

Although not depicted in detail, other bus registers include the following: a RDYBUS\_TEMPLATE\_CTL register 178 (FIG. 6D), which maintains the control information for the ready bus and the Fbus controllers, for example, it enables the ready bus sequencer; a RDYBUS\_SYNCH\_COUNT\_DEFAULT register 180 (FIG. 6E), which specifies the program cycle rate of the ready bus sequencer; and an FP\_FASTPORT\_CTL register 182 (FIG. 6F), which specifies how many Fbus clock

cycles the RSM 166 must wait between the last data transfer and the next sampling of fast receive status, as will be described.

Referring to FIG. 7A, the MAC device 30 provides transmit status flags 200 and receive status flags 202 that indicate whether the amount of data in an associated transmit FIFO 204 or receive FIFO 206 has reached a certain threshold level. The ready bus sequencer 160 periodically polls the ready flags (after selecting either the receive ready flags 202 or the transmit ready flags 200 via a flag select 208) and places them into appropriate ones of the CSRs 138 by transferring the flag data over ready bus data lines 209. In this embodiment, the ready bus includes 8 data lines for transferring flag data from each port to the Fbus interface unit 124. The CSRs in which the flag data are written are defined as RCV\_RDY\_HI/LO registers 210 for receive ready flags and XMIT\_RDY\_HI/LO registers 212 for transmit ready flags, if the ready bus sequencer 160 is programmed to execute receive and transmit ready flag read instructions, respectively.

When the ready bus sequencer is programmed with an appropriate instruction directing it to interrogate MAC receive ready flags, it reads the receive ready flags from the MAC device or devices specified in the instruction and places the flags into RCV\_RDY\_HI register 210a and a RCV\_RDY\_LO register 210b, collectively, RCV\_RDY registers 210. Each bit in these registers corresponds to a different device port on the I/O bus.

Also, and as shown in FIG. 7B, the bus interface unit 124 also supports two fast port receive ready flag pins FAST\_RX1 214a and FAST\_RX2 214b for the two fast ports of the fast MAC device 31. These fast port receive ready flag pins are read by the RSM 166 directly and placed into an RCV\_RDY\_CNT register 216.

The RCV\_RDY\_CNT register 216 is one of several used by the receive scheduler to determine how to issue a receive request. It also indicates whether a flow control request is issued.

Referring to FIG. 8A, the format of the RCV\_RDY\_CNT register 216 is as follows: bits 31:28 are defined as a reserved field 216a; bit 27 is defined as a ready bus master field 216b and is used to indicate whether the ready bus 128 is configured as a master or slave; a field corresponding to bit 26 216c provides flow control information; bits 25 and 24 correspond to FRDY2 field 216d and FRDY1 field 216e, respectively. The FRDY2 216d and FRDY1 216e are used to store the values of the FAST\_RX2 pin 214b and FAST\_RX1 pin 214a, respectively, both of which are sampled by the RSM 166 each Fbus clock cycle; bits 23:16 correspond to a reserved field 216f; a receive request count field (bits 15:8) 216g specifies a receive request count, which is incremented after the RSM 166 completes a receive request and data is available in the RFIFO 136; a receive ready count field (bits 7:0) 216h specifies a receive ready count, an 8-bit counter that is incremented each time the ready bus sequencer 160 writes the ready bus registers RCV\_RDY\_CNT register 216, the RCV\_RDY\_LO register 210b and RCV\_RDY\_HI register 210a to the receive scheduler read transfer registers.

There are two techniques for reading the ready bus registers: “autopush” and polling. The autopush instruction may be executed by the ready bus sequencer 160 during a receive process (rxautopush) or a transmit process (txautopush). Polling requires that a microengine thread periodically issue read references to the I/O bus interface 28.

The rxautopush operation performs several functions. It increments the receive ready count in the RCV\_RDY\_CNT register 216. If enabled by the RCV\_RDY\_CTL register 174, it automatically writes the RCV\_RDY\_CNT 216, the RCV\_RDY\_LO and RCV\_RDY\_HI registers 210b, 210a to the receive scheduler read

transfer registers and signals to the receive scheduler thread 92 (via a context event signal) when the rxautopush operation is complete.

The ready bus sequencer 160 polls the MAC FIFO status flags periodically and asynchronously to other events occurring in the processor 12. Ideally, the rate at which the MAC FIFO ready flags are polled is greater than the maximum rate at which the data is arriving at the MAC ports. Thus, it is necessary for the receive scheduler thread 92 to determine whether the MAC FIFO ready flags read by the ready bus sequencer 160 are new, or whether they have been read already. The rxautopush instruction increments the receive ready count in the RCV\_RDY\_CNT register 216 each time the instruction executes. The RCV\_RDY\_CNT register 216 can be used by the receive scheduler thread 92 to determine whether the state of specific flags have to be evaluated or whether they can be ignored because receive requests have been issued and the port is currently being serviced. For example, if the FIFO threshold for a Gigabit Ethernet port is set so that the receive ready flags are asserted when 64 bytes of data are in the MAC receive FIFO 206, then the state of the flags does not change until the next 64 bytes arrive 5120 ns later. If the ready bus sequencer 160 is programmed to collect the flags four times each 5120 ns period, the next three sets of ready flags that are to be collected by the ready bus sequence 160 can be ignored.

When the receive ready count is used to monitor the freshness of the receive ready flags, there is a possibility that the receive ready flags will be ignored when they are providing new status. For a more accurate determination of ready flag freshness, the receive request count may be used. Each time a receive request is completed and the receive control information is pushed onto the RCV\_CNTL register 232, the the RSM 166 increments the receive request count. The count is recorded in the RCV\_RDY\_CNT register the first time the ready bus sequencer executes an rxrdy instruction for each program loop. The receive scheduler thread 92 can use this count to

track how many requests the receive state machine has completed. As the receive scheduler thread issues commands, it can maintain a list of the receive requests it submits and the ports associated with each such request.

Referring to FIGS. 8B and 8C, the registers RCV\_RDY\_HI 210a and  
5 RCV\_RDY\_LO 210b have a flag bit 217a, 217b, respectively, corresponding to each port.

Referring to FIG. 9, the receive scheduler thread 92 performs its tasks as quickly as possible to ensure that the RSM 166 is always busy, that is, that there is always a receive request waiting to be processed by the RSM 166. Several tasks  
10 performed by the receive scheduler 92 are as follows. The receive scheduler 92 determines which ports need to be serviced by reading the RCV\_RDY\_HI, RCV\_RDY\_LO and RCV\_RDY\_CNT registers 210a, 210b and 216, respectively. The receive scheduler 92 also determines which receive ready flags are new and which are old using either the receive request count or the receive ready count in the  
15 RCV\_RDY\_CNT register, as described above. It tracks the thread processing status of the other microengine threads by reading thread done status CSRs 240. The receive scheduler thread 92 initiates transfers across the Fbus 132 via the ready bus, while the receive state machine 166 performs the actual read transfer on the Fbus 132. The receive scheduler 92 interfaces to the receive state machine 166 through two FBI CSRs  
20 138: an RCV\_REQ register 230 and an RCV\_CNTL register 232. The RCV\_REQ register 230 instructs the receive state machine on how to receive data from the Fbus 132.

Still referring to FIG. 9, a process of initiating an Fbus receive transfer is shown. Having received ready status information from the RCV\_RDY\_HI/LO registers  
25 210a, 210b as well as thread availability from the thread done register 240 (transaction "1", as indicated by the arrow labeled 1), the receive scheduler thread 92 determines if

there is room in the RCV\_REQ FIFO 230 for another receive request. If it determines that RCV\_REQ FIFO 230 has room to receive a request, the receive scheduler thread 92 writes a receive request by pushing data into the RCV\_REQ FIFO 230 (transaction 2). The RSM 166 processes the request in the RCV\_REQ FIFO 230 (transaction 3). The

5 RSM 166 responds to the request by moving the requested data into the RFIFO 136 (transaction 4), writing associated control information to the RCV\_CTL FIFO 232 (transaction 5) and generating a start\_receive signal event to the receive processing thread 96 specified in the receive request (transaction 6). The RFIFO 136 includes 16 elements 241, each element for storing a 64 byte segment of data referred to herein as a

10 MAC packet ("MPKT"). The RSM 166 reads packets from the MAC ports in fragments equal in size to one or two RFIFO elements, that is, MPKTs. The specified receive processing thread 96 responds to the signal event by reading the control information from the RCV\_CTL register 232 (transaction 7). It uses the control information to determine, among other pieces of information, where the data is located

15 in the RFIFO 136. The receive processing thread 96 reads the data from the RFIFO 136 on quadword boundaries into its read transfer registers or moves the data directly into the SDRAM (transaction 8).

The RCV\_REQ register 230 is used to initiate a receive transfer on the Fbus and is mapped to a two-entry FIFO that is written by the microengines. The I/O bus

20 interface provides signals (not shown) to the receive scheduler thread indicating that the RCV\_REQ FIFO 236 has room available for another receive request and that the last issued receive request has been stored in the RCV\_REQ register 230.

Referring to FIG. 10A, the RCV\_REQ FIFO 230 includes two entries 231. The format of each entry 231 is as follows. The first two bits correspond to a reserved

25 field 230a. Bit 29 is an FA field 230b for specifying the maximum number of Fbus accesses to be performed for this request. A THSG field (bits 28:27) 230c is a two-bit

thread message field that allows the scheduler thread to pass a message to the assigned receive thread through the ready state machine, which copies this message to the RCV\_CNTL register. An SL field 230d (bit 26) is used in cases where status information is transferred following the EOP MPKT. It indicates whether two or one 32-bit bus accesses are required in a 32-bit Fbus configuration. An E1 field 230e (bits 21:18) and an E2 field (bits 25:22) 230f specify the RFIFO element to receive the transferred data. If only 1 MPKT is received, it is placed in the element indicated by the E1 field. If two MPKTs are received, then the second MPKT is placed in the RFIFO element indicated by the E2 field. An FS field (bits 17:16) 230g specifies use of a fast or slow port mode, that is, whether the request is directed to a fast or slow port. The fast port mode setting signifies to the RSM that a sequence number is to be associated with the request and that it will be handling speculative requests, which will be discussed in further detail later. An NFE field (bit 15) 230h specifies the number of RFIFO elements to be filled (i.e., one or two elements). The IGFR field (bit 13) 230i is used only if fast port mode is selected and indicates to the RSM that it should process the request regardless of the status of the fast ready flag pins. An SIGRS field (bit 11) 230j, if set, indicates that the receive scheduler be signaled upon completion of the receive request. A TID field (bits 10:6) 230k specifies the receive thread to be notified or signaled after the receive request is processed. Therefore, if bit 11 is set, the RCV\_REQ entry must be read twice, once by the receive thread and once by the receive scheduler thread, before it can be removed from the RCV\_REQ FIFO. An RM field (bits 5:3) 230l specifies the ID of the MAC device that has been selected by the receive scheduler. Lastly, an RP field (bits 2:0) 230m specifies which port of the MAC device specified in the RM field 230l has been selected.

The RSM 166 reads the RCV\_REQ register entry 231 to determine how it should receive data from the Fbus 132, that is, how the signaling should be performed

on the Fbus, where the data should be placed in the RFIFO and which microengine thread should be signaled once the data is received. The RSM 166 looks for a valid receive request in the RCV\_REQ FIFO 230. It selects the MAC device identified in the RM field and selects the specified port within the MAC by asserting the appropriate control signals. It then begins receiving data from the MAC device on the Fbus data lines. The receive state machine always attempts to read either eight or nine quadwords of data from the MAC device on the Fbus as specified in the receive request. If the MAC device asserts the EOP signal, the RSM 166 terminates the receive early (before eight or nine accesses are made). The RSM 166 calculates the total bytes received for each receive request and reports the value in the REC\_CNTL register 232. If EOP is received, the RSM 166 determines the number of valid bytes in the last received data cycle.

The RCV\_CNTL register 232 is mapped to a four-entry FIFO (referred to herein as RCV\_CNTL\_FIFO 232) that is written by the receive state machine and read by the microengine thread. The I/O bus interface 28 signals the assigned thread when a valid entry reaches the top of the RCV\_CNTL\_FIFO. When a microengine thread reads the RCV\_CNTL register, the data is popped off the FIFO. If the SIGRS field 230i is set in the RCV\_REQ register 230, the receive scheduler thread 92 specified in the RCV\_CNTL register 232 is signaled in addition to the thread specified in TID field 230k. In this case, the data in the RCV\_CNTL register 232 is read twice before the receive request data is retired from the RCV\_CTL\_FIFO 232 and the next thread is signaled. The receive state machine writes to the RCV\_CTL register 232 as long as the FIFO is not full. If the RCV\_CTL\_FIFO 232 is full, the receive state machine stalls and stops accepting any more receive requests.

Referring to FIG. 10B, the RCV\_CNTL\_FIFO 232 provides instruction to the signaled thread (i.e., the thread specified in TID) to process the data. As indicated

above, the RCV\_CNTL FIFO includes 4 entries 233. The format of the RCV\_CNTL FIFO entry 233 is as follows: a THMSG field (31:30) 23a includes the 2-bit message copied by the RSM from REC\_REQ register[28:27]. A MACPORT/THD field (bits 29:24) 232b specifies either the MAC port number or a receive thread ID, as will be described in further detail below. An SOP\_SEQ field (23:20) 232c is used for fast ports and indicates a packet sequence number as an SOP (start-of-packet) sequence number if the SOP was asserted during the receive data transfer and indicates an MPKT sequence number if SOP was not so asserted. An RF field 232d and RERR field 232e (bits 19 and 18, respectively) both convey receive error information. An SE field 232f (17:14) and an FE field 232g (13:10) are copies of the E2 and E1 fields, respectively, of the REC\_REQ. An EF field (bit 9) 232h specifies the number of RFIFO elements which were filled by the receive request. An SN field (bit 8) 232i is used for fast ports and indicates whether the sequence number specified in SOP\_SEQ field 232c is associated with fast port 1 or fast port 2. A VLD\_BYTES field (7:2) 232j specifies the number of valid bytes in the RFIFO element if the element contains in EOP MPKT. An EOP field (bit 1) 232k indicates that the MPKT is an EOP MPKT. An SOP field (bit 0) 232l indicates that the MPKT is an SOP MPKT.

FIG. 11 illustrates the format of the thread done registers 240 and their interaction with the receive scheduler and processing threads 92, 96, respectively, of the microengines 22. The thread done registers 240 include a first thread status register, TH\_DONE\_REG0 240a, which has 2-bit status fields 241a corresponding to each of threads 0 through 15. A second thread status register, TH\_DONE\_REG1 240b, has 2-bit status fields 241b corresponding to each of threads 16 through 23. These registers can be read and written to by the threads using a CSR instruction (or fast write instruction, described below). The receive scheduler thread can use these registers to determine which RFIFO elements are not in use. Since it is the receive scheduler thread

92 that assigns receive processing threads 96 to process the data in the RFIFO elements, and it also knows the thread processing status from the THREAD\_DONE\_REG0 and THREAD\_DONE\_REG1 registers 240a, 240b, it can determine which RFIFO elements are currently available.

5           The THREAD\_DONE CSRs 240 support a two-bit message for each microengine thread. The assigned receive thread may write a two-bit message to this register to indicate that it has completed its task. Each time a message is written to the THREAD\_DONE register, the current message is logically ORed with the new message. The bit values in the THREAD\_DONE registers are cleared by writing a “1”,  
10 so the scheduler may clear the messages by writing the data read back to the THREAD\_DONE register. The definition of the 2-bit status field is determined in software. An example of four message types is illustrated in TABLE 1 below.

2-BIT MESSAGE	DEFINITION
00	Busy.
01	Idle, processing complete.
10	Not busy, but waiting to finish processing of entire packet.
11	Idle, processing complete for an EOP MPKT.

TABLE 1

The assigned receive processing threads write their status to the THREAD\_DONE register whenever the status changes. For example, a thread may immediately write 00 to the THREAD\_DONE register after the receive state machine signals the assigned  
20 thread. When the receive scheduler thread reads the THREAD\_DONE register, it can

look at the returned value to determine the status of each thread and then update its thread/port assignment list.

The microengine supports a fast\_wr instruction that improves performance when writing to a subset of CSR registers. The fast\_wr instruction does not use the push or pull engines. Rather, it uses logic that services the instruction as soon as the write request is issued to the FBI CSR. The instruction thus eliminates the need for the pull engine to read data from a microengine transfer register when it processes the command. The meaning of the 10-bit immediate data for some of the CSRs is shown below.

CSR	10-BIT IMMEDIATE DATA
INTER_THD_SIG	Thread number of the thread that is to be signaled.
THREAD_DONE	A 2-bit message that is shifted into a position relative to the thread that is writing the message.
THREAD_DONE_INCR1 THREAD_DONE_INCR2	Same as THREAD_DONE except that either the enqueue_seq1 or enqueue_seq2 is also incremented.
INCR_ENQ_NUM1 INCR_ENQ_NUM2	Write a one to increment the enqueue sequence number by one.

TABLE 2

It will be appreciated that the receive process as described herein assumes that no packet exemptions occurred, that is, that the threads are able to handle the packet processing without assistance from the core processor. Further, the receive process as described also assumes the availability of FIFO space. It will be appreciated that the various state machines must determine if there is room available in a FIFO, e.g., the RFIFO, prior to writing new entries to that FIFO. If a particular FIFO is full, the state

machine will wait until the appropriate number of entries has been retired from that FIFO.

#### Scheduling of Received Network Data

In one embodiment, one of the microengines 22a-f is designated as a receive scheduler 92 and three microengines 22a-f are designated to be receive threads 96. Because there are four threads per microengine, there are twelve receive threads total. The receive scheduler 92 assigns a ready port to an available receive thread, and causes transfer of packet data into the network processor. The receive threads then analyze and store sections of incoming packet data (MPKTs). Since the receive threads can operate simultaneously, multiple portions of a MPKT can be processed by receive threads 96 at the same time.

FIG. 12 shows a graphical overview of the scheduling process showing an exemplary receive scheduler and three receive threads. The scheduling process will be described in more detail with respect to Figures 13 and 14. In Figure 12, the receive scheduler initiates three receive requests one after another (301). Each receive request initiates a bus transfer that wakes up an available thread (302). Thread X 303 is the first thread to wake up. It retrieves control information and saves state, as will be described in more detail later, and then continues to process its network data. Passing state information early on in the thread processing allows the network processor to maintain a high data rate since the threads can continue to process their packet section at their own slower pace.

Thread Y 304 wakes up from the second bus transfer. Similar to thread X, thread Y retrieves control information, saves state, then processes its network data. Thread Z 305 behaves similarly. If threads X, Y, and Z are in different microengines, then each thread may proceed without context switching relative to each other.

FIG. 13 shows a flowchart of the interaction of processes between the scheduler thread 92 and an available receive thread 96. At box 401, the receive scheduler 92 obtains port read status of the ports and thread available status to determine which ports and threads are available. The receive scheduler 92 selects one of the available threads (box 401). The receive scheduler 92 also selects a ready port (box 402). A port is ready if it has data in its receive buffer that meets some criteria (e.g., received number of bytes over a threshold, received a complete packet, etc). The receive scheduler 92 assembles the receive request, inserting the port number and ID of the receive thread to wake up (box 403).

The receive scheduler 92 increments a mailbox slot id (called MSEQ, for MPKT sequence) (box 404). The mailbox can be implemented in a register or fast memory. In one embodiment, the mailbox has sixteen slots and the receive scheduler 92 increments MSEQ then truncates it to four bits, ensuring that MSEQ will count from 0-15 then wrap around to 0 and repeat on successive iterations of the loop. The receive scheduler 92 inserts the MSEQ into an unused field of the receive request (box 405). The bus interface gets the request (box 406), then starts a hardware sequencer that selects the device/port and transfers data across the bus into a receive FIFO and wakes up the receive thread (box 410).

The receive thread, upon waking up (box 420), reads the receive control register and extracts the MSEQ value from it (box 421). The receive thread then restores information about the packet (box 422), using a test and clear operation on bit 31 of the first word in the MSEQ. Bit 31 set indicates valid information. Bit 31 is set to 1 by a previous thread, or by initialization. (Note that initially the first slot of MSEQ has bit 31 set to enable the very first thread to break out of the test and clear loop.) The receive thread then modifies the information as needed, and stores it to the mailbox slot MSEQ+1 (box 423), for the next thread to retrieve. The receive thread continues to

processes its network data (box 424), writes thread done and either swaps out context or waits on a new start receive (box 425).

FIG. 14 shows a flowchart showing the interaction of processing by multiple receive threads. The packet data is received on one or more ports that are connected to MAC devices, that are in turn connected to networks. The data is received as a series of MPKTs. For every receive request transaction the port bus interface transfers a section of the packet from a MAC device port to a receive FIFO in the network processor. The bus interface wakes up a designated receive thread. The receive thread then reads a control register (box 501) to obtain information about the bus transaction, i.e., which port, where in the receive FIFO the data is stored, size of data, type of port, message and other information. In the receive control register, there can be a cancel message that indicates that the port was not ready. If there is no cancel message, the receive thread restores the packet address from the mailbox slot indexed by MSEQ (box 502), increments it by the buffer size of the data (box 504), and saves it to the mailbox slot indexed by MSEQ+1 (box 505). After saving the updated address, the thread copies the data from RFIFO to DRAM (boxes 506, 520). The receive thread writes a thread done register to indicate it is available for re-assignment.

If there is a cancel (box 503), the receive thread does not increment the address, but merely restores it from mailbox indexed by MSEQ and saves it to mailbox indexed by MSEQ+1. In this case no data is copied to DRAM. The receive thread writes a thread done register to indicate it is available for re-assignment.

The sequence repeats, with the receive scheduler making requests and assigning mailbox slots, and receive threads processing the data. In one embodiment, the mailbox slot ids are assigned incrementally by the receive scheduler 92 and passed through the bus interface to the receive threads 96. They provide that successive receives from the same port will copy data to sequential buffers in DRAM.

The mailbox entry can be used to carry a variety of descriptive information used to process the packet. A few examples include: 1) it can carry a buffer freelist identifier that identifies which memory allocation freelist was used for the packet. This information is passed through the mailbox and ultimately to the transmit threads, which  
5 can return the buffer to the correct freelist; 2) it can carry an offset into the packet buffer, that will later be used by transmit threads to locate the beginning of the packet in the memory buffer. This is useful when packets are modified as part of a protocol conversion, resulting in a different packet size at transmit.

In one embodiment, the mailbox entry includes a relative buffer address that  
10 includes a sixteen bit value. The relative buffer address (or a scaled offset thereof) is added to a SRAM base address to provide a pointer to a packet link descriptor location in SRAM. Similarly, the relative buffer address (or a scaled offset thereof) is added to a SDRAM base address to provide a pointer to packet data location in SDRAM. The mailbox may also include such as information as an element count indicating the  
15 number of 64 byte segments of data in a data packet, and the SOP sequence number. The thread that processes an EOP MPKT writes the packet link descriptor to the right place. This may be to the transmit queue as previously described.

Thus, a method of processing network data in a network processor includes using a receive scheduler to schedule multiple threads. The scheduler allows greater  
20 flexibility by governing the rate at which each port or all ports receive data, causing selected discards. However, the specific arrangements and methods described herein are merely illustrative of the principles of this invention. Numerous modifications in form and detail may be made without departing from the scope of the described invention. Although this invention has been shown in relation to a particular  
25 embodiment, it should not be considered so limited. Rather, the described invention is limited only by the scope of the appended claims.

## CLAIMS

WHAT IS CLAIMED IS:

- 1           1. A method of processing network data in a network processor comprising:  
2                 scheduling a first thread to process a first incoming block of data; and  
3                 scheduling a second thread to process a second incoming block of data prior  
4                 to the first thread completing.
- 1           2. The method of claim 1, wherein the first incoming block of data and the second  
2 incoming block of data are from a common data packet.
- 1           3. The method of claim 2 further comprising:  
2                 saving state information by the first thread; and  
3                 retrieving the state information by the second thread.
- 1           4. The method of claim 3, wherein the state information includes a pointer into a  
2 memory indicating where to move the first and second incoming blocks of data.
- 1           5. The method of claim 4 further comprising:  
2                 storing data to memory in a sequential ordering based on the state  
3                 information.
- 1           6. The method of claim 5 further comprising:

2 providing the state information to transmit circuitry.

1 7. A method of processing a data packet received over a network comprising:  
2 processing a first portion of the data packet using a first thread; and  
3 simultaneously processing a second portion of the data packet using a second  
4 thread.

1 8. The method of claim 7 wherein the first thread and the second thread do not  
2 time share processing with one another.

1 9. The method of claim 8 wherein the first thread and the second thread operate  
2 out of different microengines.

1 10. The method of claim 7 wherein the first thread and the second thread time  
2 share processing with one another.

1 11. The method of claim 10 wherein the first thread and the second thread operate  
2 out of a common microengine.

1 12. The method of claim 7 further comprising:  
2 simultaneously with processing the first portion and the second portion of  
3 the data packet, processing a third portion of the data packet using a  
4 third thread.

1           13. The method of claim 12 wherein the first thread, the second thread, and the  
2           third thread run the same code.

1           14. The method of claim 13 wherein the first thread, the second thread, and the  
2           third thread do not time share processing with one another.

1           15. An article comprising a computer-readable medium which store computer-  
2           executable instructions for receiving data from a plurality of ports, the instructions  
3           causing a computer to:

4                     process a first portion of a data packet using a first thread; and  
5                     process a second portion of the data packet using a second thread, wherein  
6                     there is no time sharing between the first thread and the second thread.

1           16. The article of claim 15, the article further comprises instructions to:  
2                     save state information of the first thread; and  
3                     restore the state information by the second thread.

1           17. The article of claim 16, the article further comprises instructions to:  
2                     provide the state information to transmit circuitry when an end of packet is  
3                     detected by a subsequent thread.

## ABSTRACT OF THE DISCLOSURE

A method of processing network data in a network processor includes scheduling a first thread to process a first incoming block of data and scheduling a second thread to process a second incoming block of data. The first and second thread  
5 can operate simultaneously.




FIG. 1

22a




FIG. 2










FIG. 6A




FIG. 6B




FIG. 6C




FIG. 6D




FIG. 6E




FIG. 6F




FIG. 7A




FIG. 7B




FIG. 8A




FIG. 8B




FIG. 8C



231	230a	230b	230c	230d	230e	230f	230g	230h	230i	230j	230k	230l	230m	230
RES	FA		JL	EZ	EI	FS	NFE	IGFR	STCAS	TID	RM	RY		C
														1

FIG. 10A

233	232a	232b	232c	232d	232e	232f	232g	232h	232i	232j	232k	232l	232
THMSG	MACRO/TID	SOP SEQ#	AF	REAR	SE	PE	EFOR	SINOR	Valid bytes	EOP	SOP		0
													1
													2
													3

FIG. 10B

09626535-072200

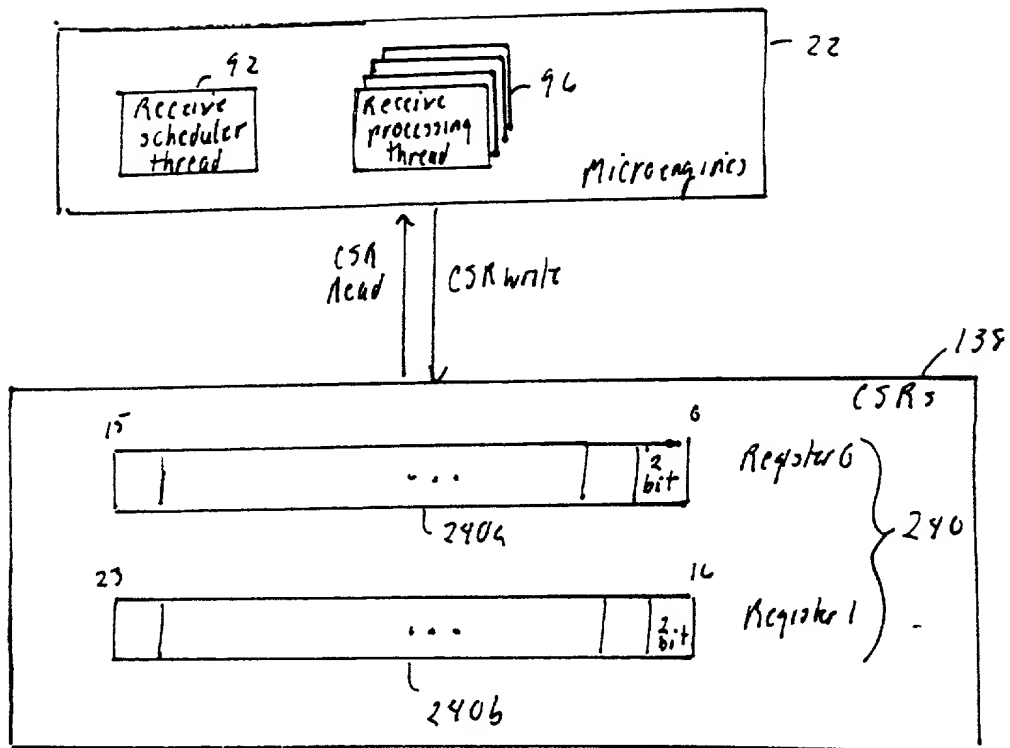


FIG. 11

FIG 12

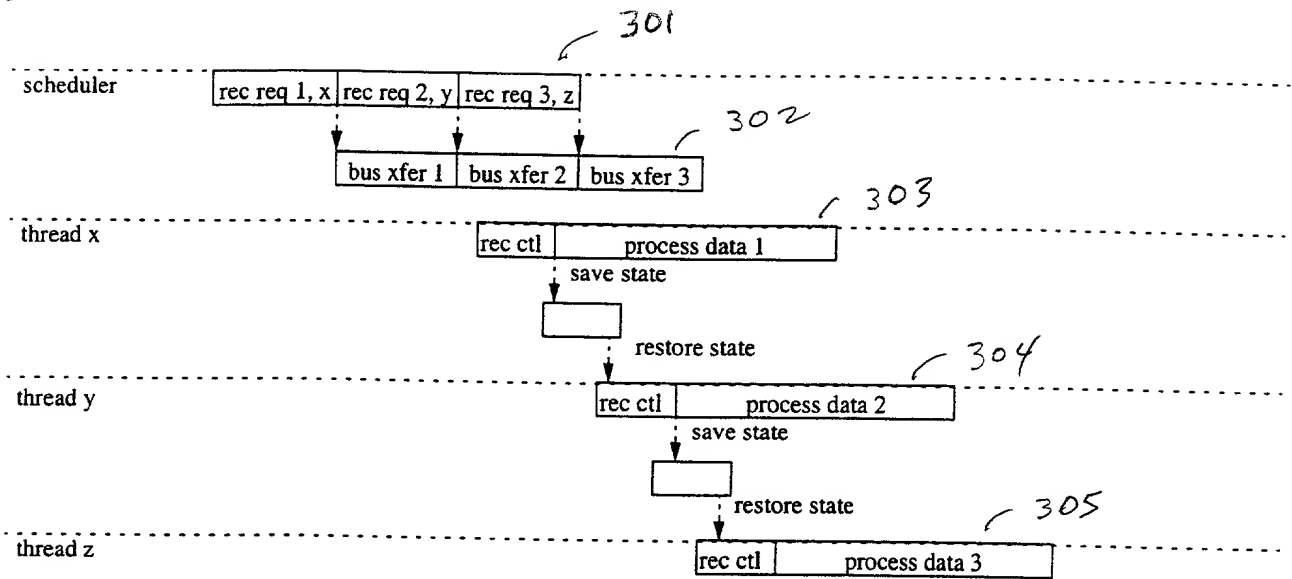


FIG 13

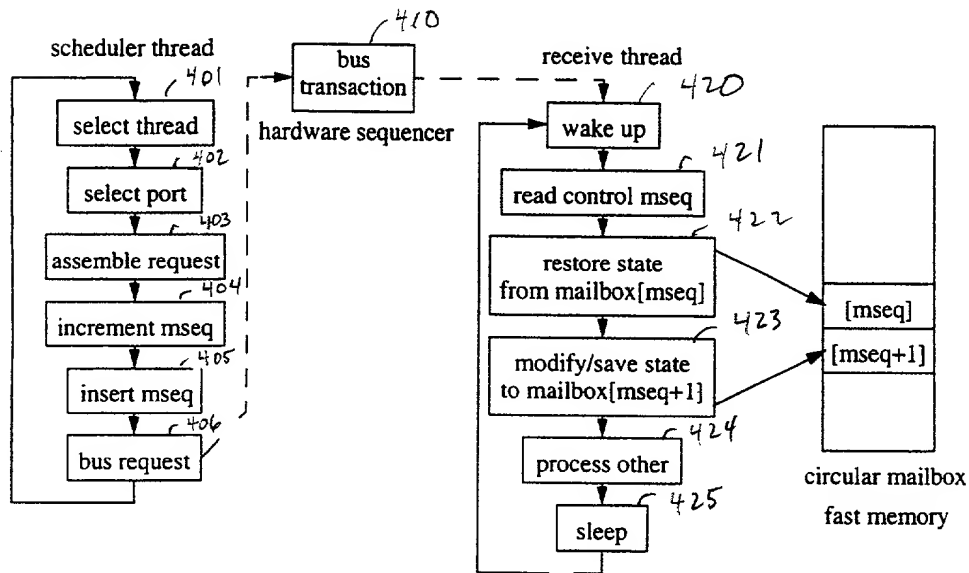



FIG. 14



DECLARATION AND POWER OF ATTORNEY FOR PATENT APPLICATION  
(CONTINUATION-IN-PART)

As a below named inventor, I hereby declare that:

My residence, post office address and citizenship are as stated below, next to my name.

I believe I am the original, first, and sole inventor (if only one name is listed below) or an original, first, and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled Multi-Threaded Scheduled Receive for Fast Network Port Data

the specification of which

XXX is attached hereto.  
\_\_\_\_\_ was filed on \_\_\_\_\_ as  
United States Application Number \_\_\_\_\_  
or PCT International Application Number \_\_\_\_\_  
and was amended on \_\_\_\_\_  
(if applicable)

I hereby state that I have reviewed and understand the contents of the above-identified specification, including the claims, as amended by any amendment referred to above.

I acknowledge the duty to disclose all information known to me to be material to patentability as defined in Title 37, Code of Federal Regulations, Section 1.56.

I hereby claim foreign priority benefits under Title 35, United States Code, Section 119(a)-(d), of any foreign application(s) for patent or inventor's certificate listed below and have also identified below any foreign application for patent or inventor's certificate having a filing date before that of the application on which priority is claimed:

<u>Prior Foreign Application(s)</u>			<u>Priority Claimed</u>	
<u>(Number)</u>	<u>(Country)</u>	<u>(Day/Month/Year Filed)</u>	<u>Yes</u>	<u>No</u>
_____	_____	_____	<u>Yes</u>	<u>No</u>
_____	_____	_____	<u>Yes</u>	<u>No</u>
_____	_____	_____	<u>Yes</u>	<u>No</u>

I hereby claim the benefit under title 35, United States Code, Section 119(e) of any United States provisional application(s) listed below:

_____	_____
Application Number	Filing Date
_____	_____
Application Number	Filing Date

I hereby claim the benefit under Title 35, United States Code, Section 120 of any United States application(s) listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States application in the manner provided by the first paragraph of Title 35, United States Code, Section 112, I acknowledge the duty to disclose all information known to me to be material to patentability as defined in Title 37, Code of Federal Regulations, Section 1.56 which became available between the filing date of the prior application and the national or PCT international filing date of this application:

<u>09/475,614</u>	<u>12/30/99</u>	<u>pending</u>
Application Number	Filing Date	Status -- patented, pending, abandoned

Application Number	Filing Date	Status -- patented, pending, abandoned
--------------------	-------------	---

I hereby appoint the persons listed on Appendix A hereto (which is incorporated by reference and a part of this document) as my respective patent attorneys and patent agents, with full power of substitution and revocation, to prosecute this application and to transact all business in the Patent and Trademark Office connected herewith.

Send correspondence to Charles K. Young, BLAKELY, SOKOLOFF, TAYLOR &  
(Name of Attorney or Agent)  
ZAFMAN LLP, 12400 Wilshire Boulevard 7th Floor, Los Angeles, California 90025 and direct  
telephone calls to Maria McCormack Sobrino, (408) 720-8300.  
(Name of Attorney or Agent)

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

Full Name of Sole/First Inventor Donald F. Hooper

Inventor's Signature \_\_\_\_\_ Date \_\_\_\_\_

Residence Shrewsbury, MA Citizenship USA  
(City, State) (Country)

Post Office Address 19 Main Circle  
Shrewsbury, MA 01545

Full Name of Second/Joint Inventor Matthew J. Adiletta

Inventor's Signature \_\_\_\_\_ Date \_\_\_\_\_

Residence Worcester, MA Citizenship USA  
(City, State) (Country)

Post Office Address 20 Monticello Drive  
Worcester, MA 01603

00220"9392950

Full Name of Third/Joint Inventor Gilbert M. Wolrich

Inventor's Signature \_\_\_\_\_ Date \_\_\_\_\_

Residence Framingham, MA \_\_\_\_\_ Citizenship USA \_\_\_\_\_  
(City, State) (Country)

Post Office Address 4 Cider Mill Road \_\_\_\_\_  
Framingham, MA 01701 \_\_\_\_\_

Full Name of Fourth/Joint Inventor Debra Bernstein

Inventor's Signature \_\_\_\_\_ Date \_\_\_\_\_

Residence Sudbury, MA \_\_\_\_\_ Citizenship USA \_\_\_\_\_  
(City, State) (Country)

Post Office Address 443 Peakham Road \_\_\_\_\_  
Sudbury, MA 01776 \_\_\_\_\_

Full Name of Fifth/Joint Inventor \_\_\_\_\_

Inventor's Signature \_\_\_\_\_ Date \_\_\_\_\_

Residence \_\_\_\_\_ Citizenship \_\_\_\_\_  
(City, State) (Country)

Post Office Address \_\_\_\_\_  
\_\_\_\_\_

Full Name of Sixth/Joint Inventor \_\_\_\_\_

Inventor's Signature \_\_\_\_\_ Date \_\_\_\_\_

Residence \_\_\_\_\_ Citizenship \_\_\_\_\_  
(City, State) (Country)

Post Office Address \_\_\_\_\_  
\_\_\_\_\_

Full Name of Seventh/Joint Inventor \_\_\_\_\_

Inventor's Signature \_\_\_\_\_ Date \_\_\_\_\_

Residence \_\_\_\_\_ Citizenship \_\_\_\_\_  
(City, State) (Country)

Post Office Address \_\_\_\_\_  
\_\_\_\_\_

## APPENDIX A

William E. Alford, Reg. No. 37,764; Farzad E. Amini, Reg. No. P42,261; Aloysius T. C. AuYeung, Reg. No. 35,432; William Thomas Babbitt, Reg. No. 39,591; Carol F. Barry, Reg. No. 41,600; Jordan Michael Becker, Reg. No. 39,602; Lisa N. Benado, Reg. No. 39,995; Bradley J. Bereznak, Reg. No. 33,474; Michael A. Bernadicou, Reg. No. 35,934; Roger W. Blakely, Jr., Reg. No. 25,831; R. Alan Burnett, Reg. No. 46,149; Gregory D. Caldwell, Reg. No. 39,926; Andrew C. Chen, Reg. No. 43,544; Thomas M. Coester, Reg. No. 39,637; Donna Jo Coningsby, Reg. No. 41,684; Florin Corie, Reg. No. 46,244; Dennis M. deGuzman, Reg. No. 41,702; Stephen M. De Klerk, Reg. No. P46,503; Michael Anthony DeSanctis, Reg. No. 39,957; Daniel M. De Vos, Reg. No. 37,813; Robert Andrew Diehl, Reg. No. 40,992; Sanjeet Dutta, Reg. No. P46,145; Matthew C. Fagan, Reg. No. 37,542; Tarek N. Fahmi, Reg. No. 41,402; George Fountain, Reg. No. 37,374; Paramita Ghosh, Reg. No. 42,806; James Y. Go, Reg. No. 40,621; Libby N. Ho, Reg. No. P46,774; James A. Henry, Reg. No. 41,064; Willmore F. Holbrow III, Reg. No. P41,845; Sheryl Sue Holloway, Reg. No. 37,850; George W. Hoover II, Reg. No. 32,992; Eric S. Hyman, Reg. No. 30,139; William W. Kidd, Reg. No. 31,772; Sang Hui Kim, Reg. No. 40,450; Walter T. Kim, Reg. No. 42,731; Eric T. King, Reg. No. 44,188; Erica W. Kuo, Reg. No. 42,775; George Brian Leavell, Reg. No. 45,436; Kurt P. Leyendecker, Reg. No. 42,799; Gordon R. Lindeen III, Reg. No. 33,192; Jan Carol Little, Reg. No. 41,181; Joseph Lutz, Reg. No. 43,765; Michael J. Mallie, Reg. No. 36,591; Andre L. Marais, under 37 C.F.R. § 10.9(b); Paul A. Mendonsa, Reg. No. 42,879; Clive D. Menezes, Reg. No. 45,493; Chun M. Ng, Reg. No. 36,878; Thien T. Nguyen, Reg. No. 43,835; Thinh V. Nguyen, Reg. No. 42,034; Dennis A. Nicholls, Reg. No. 42,036; Daniel E. Ovanezian, Reg. No. 41,236; Kenneth B. Paley, Reg. No. 38,989; Marina Portnova, Reg. No. P45,750; William F. Ryann, Reg. No. 44,313; James H. Salter, Reg. No. 35,668; William W. Schaal, Reg. No. 39,018; James C. Scheller, Reg. No. 31,195; Jeffrey Sam Smith, Reg. No. 39,377; Maria McCormack Sobrino, Reg. No. 31,639; Stanley W. Sokoloff, Reg. No. 25,128; Judith A. Szepesi, Reg. No. 39,393; Vincent P. Tassinari, Reg. No. 42,179; Edwin H. Taylor, Reg. No. 25,129; John F. Travis, Reg. No. 43,203; Joseph A. Twarowski, Reg. No. 42,191; Tom Van Zandt, Reg. No. 43,219; Lester J. Vincent, Reg. No. 31,460; Glenn E. Von Tersch, Reg. No. 41,364; John Patrick Ward, Reg. No. 40,216; Mark L. Watson, Reg. No. P46,322; Thomas C. Webster, Reg. No. P46,154; Steven D. Yates, Reg. No. 42,242; and Norman Zafman, Reg. No. 26,250; my patent attorneys, and Firasat Ali, Reg. No. 45,715; and Justin M. Dillon, Reg. No. 42,486; my patent agents, of BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN LLP, with offices located at 12400 Wilshire Boulevard, 7th Floor, Los Angeles, California 90025, telephone (310) 207-3800, and Alan K. Aldous, Reg. No. 31,905; Edward R. Brake, Reg. No. 37,784; Ben Burge, Reg. No. 42,372; Jeffrey S. Draeger, Reg. No. 41,000; Cynthia Thomas Faatz, Reg. No. 39,973; John N. Greaves, Reg. No. 40,362; Seth Z. Kalson, Reg. No. 40,670; David J. Kaplan, Reg. No. 41,105; Peter Lam, Reg. No. 44,855; Charles A. Mirho, Reg. No. 41,199; Leo V. Novakoski, Reg. No. 37,198; Thomas C. Reynolds, Reg. No. 32,488; Kenneth M. Seddon, Reg. No. 43,105; Mark Seeley, Reg. No. 32,299; Steven P. Skabrat, Reg. No. 36,279; Howard A. Skaist, Reg. No. 36,008; Gene I. Su, Reg. No. 45,140; Calvin E. Wells, Reg. No. P43,256; Raymond J. Werner, Reg. No. 34,752; Robert G. Winkle, Reg. No. 37,474; and Charles K. Young, Reg. No. 39,435; my patent attorneys, of INTEL CORPORATION; and James R. Thein, Reg. No. 31,710, my patent attorney with full power of substitution and revocation, to prosecute this application and to transact all business in the Patent and Trademark Office connected herewith.

## APPENDIX B

### Title 37, Code of Federal Regulations, Section 1.56 Duty to Disclose Information Material to Patentability

(a) A patent by its very nature is affected with a public interest. The public interest is best served, and the most effective patent examination occurs when, at the time an application is being examined, the Office is aware of and evaluates the teachings of all information material to patentability. Each individual associated with the filing and prosecution of a patent application has a duty of candor and good faith in dealing with the Office, which includes a duty to disclose to the Office all information known to that individual to be material to patentability as defined in this section. The duty to disclosure information exists with respect to each pending claim until the claim is cancelled or withdrawn from consideration, or the application becomes abandoned. Information material to the patentability of a claim that is cancelled or withdrawn from consideration need not be submitted if the information is not material to the patentability of any claim remaining under consideration in the application. There is no duty to submit information which is not material to the patentability of any existing claim. The duty to disclose all information known to be material to patentability is deemed to be satisfied if all information known to be material to patentability of any claim issued in a patent was cited by the Office or submitted to the Office in the manner prescribed by §§1.97(b)-(d) and 1.98. However, no patent will be granted on an application in connection with which fraud on the Office was practiced or attempted or the duty of disclosure was violated through bad faith or intentional misconduct. The Office encourages applicants to carefully examine:

- (1) Prior art cited in search reports of a foreign patent office in a counterpart application, and
- (2) The closest information over which individuals associated with the filing or prosecution of a patent application believe any pending claim patentably defines, to make sure that any material information contained therein is disclosed to the Office.

(b) Under this section, information is material to patentability when it is not cumulative to information already of record or being made or record in the application, and

- (1) It establishes, by itself or in combination with other information, a prima facie case of unpatentability of a claim; or
- (2) It refutes, or is inconsistent with, a position the applicant takes in:
  - (i) Opposing an argument of unpatentability relied on by the Office, or
  - (ii) Asserting an argument of patentability.

A prima facie case of unpatentability is established when the information compels a conclusion that a claim is unpatentable under the preponderance of evidence, burden-of-proof standard, giving each term in the claim its broadest reasonable construction consistent with the specification, and before any consideration is given to evidence which may be submitted in an attempt to establish a contrary conclusion of patentability.

(c) Individuals associated with the filing or prosecution of a patent application within the meaning of this section are:

- (1) Each inventor named in the application;
- (2) Each attorney or agent who prepares or prosecutes the application; and
- (3) Every other person who is substantively involved in the preparation or prosecution of the application and who is associated with the inventor, with the assignee or with anyone to whom there is an obligation to assign the application.

(d) Individuals other than the attorney, agent or inventor may comply with this section by disclosing information to the attorney, agent, or inventor.